

An Example of Bayesian Inference

Dell Zhang
16/11/2008

There's a disease that is affecting 1% of the population. We have a test for the disease that's only somewhat accurate: 80% of people with the disease will get a positive result. Unfortunately, 9.6% of people without the disease will also get a positive result. We apply the test to you and it turns out positive. What is the probability that you actually have the disease?

// Prior Probability

$$P(D=1) = 0.01 \quad P(D=0) = 0.99$$

// Likelihood

$$P(+|D=1) = 0.8 \quad P(+|D=0) = 0.096$$

$$\begin{aligned} P(+) &= P(+,D=1) + P(+,D=0) && // \text{Partition Rule} \\ &= P(+|D=1) P(D=1) + P(+|D=0) P(D=0) && // \text{Chain Rule} \\ &= 0.8*0.01 + 0.096*0.99 \\ &\approx 0.103 \end{aligned}$$

// Posterior Probability computation using the Bayes' Rule

$$\begin{aligned} P(D=1|+) &= P(+,D=1) / P(+) && // \text{Chain Rule} \\ &= P(+|D=1) P(D=1) / P(+) && // \text{Chain Rule} \\ &= 0.8*0.01 / 0.103 \\ &\approx 0.078 = 7.8\% \end{aligned}$$

P.S.

Consider a population of 10,000 persons.

Of those, 100 will have the disease, and 80 will test positive.

That leaves 10000-100 who won't have the disease, or 9900.

The number of folks who don't have the disease but tested positive is 9.6% of 9900, or 950.

$80 / (950 + 80) = 0.078$ or 7.8%.